

NOTIFICATION PLATFORM ARCHITECTURE

RELATED APPLICATIONS

The present application claims the benefit of the provisional patent application filed on March 16, 2000, entitled "Attentional Systems and Interfaces," and assigned
5 serial no. 60/189,801 [attorney docket no. 1018.101US0].

FIELD OF THE INVENTION

This invention relates generally to unified receipt and notification of alerts generated by varied devices and applications for conveyance to a user, and more particularly to an architecture for such unified context sensing and analysis, alert receipt,
10 and notification.

BACKGROUND OF THE INVENTION

Many computer users today receive information from a number of different sources, and utilize a number of different devices in order to access this information. For example, a user may receive e-mail and instant messages over a computer, pages over a
15 pager, voice-mail over a phone, such as a cellular ("cell") or landline phone, news information over the computer, etc. This makes it difficult for the user to receive all his or her different information wherever the user happens to be.

For example, a user may be away from his or her computer, but receive an important e-mail. The user may have access only to a cell phone or a pager, however.
20 As another example, the user may be working on the computer, and have turned off the

ringer and voice-mail indicator on the phone. When an important voice-mail is left, the user has no way of receiving this information on the computer.

Moreover, many of the alerts may not be important to the user – for example, an e-mail from the user’s manager or co-worker should receive higher priority than the latest sports scores. More generally, the value of the information contained in an alert should be balanced with the costs associated with the disruption of the user by an alert. Both the costs and value may be context sensitive. Beyond notifications about communications, users are alerted with increasing numbers of services, error messages, and computerized offers for assistance.

The prior art, however, does not provide for alerts following the user, for the prioritization of the alerts, nor for considering the potentially context-sensitive value and costs associated with notifications. For these and other reasons, therefore, there is a need for the present invention.

SUMMARY OF THE INVENTION

This invention relates to the architecture for a notification platform. In one embodiment, the architecture includes a user mechanism, one or more notification sources and sinks, and a notification manager. The user mechanism is designed to both store user profile information regarding notification parameters of a user, such as the user’s default notification preferences, and to provide a user context identification and updating service. Each notification source is designed to generate notifications intended for the user, while each notification sink is designed to provide the notifications to the user. The notification manager is designed to convey the notifications generated by the sources to the sinks, based on information stored by the user mechanism and on

information provided or inferred about the urgencies of the notifications. For example, the notification manager can access or infer the context of the user (e.g., the user's current location and focus of attention). This can be done based on a consideration of multiple sources of information. Such sources of information can include the user's context profile, the user's online calendar, the time of day, events about the world, organization, system, and/or the user's activity. The notification can then determine through an analysis of the context and the urgency of the information which of the notifications should be conveyed to the user, via which of the sinks, and in which manner or modality provided by the sinks.

Embodiments of the invention provide for advantages. A user may, for example, receive e-mail alerts on a cell phone, or voice-mail on a desktop computer, as is determined appropriate by the notification manager. Thus, notifications from notification sources are passed to the notification manager, which determines whether the user should be notified. If the manager determines that the user should be notified, then the manager also determines how the user should be notified. This can be based on the information stored in the user profile, including such information as the user's preferences and current context, and notifies the appropriate notification sinks. The sinks can include, for example, a desktop computer, a cell phone, a pager, etc.

Furthermore, the methods and architecture of the notification platform generalize to any notification, including those associated with the potential provision of a service by a software component in a desktop or mobile setting. Such notifications include:

- alerts about services such as those that seek to automatically provide assistance or tips to a user working with a software application and/or automatically perform scheduling by examining e-mail at the focus of a user's attention;
- alerts that notify the user to upcoming appointments or engagements;
- 5 • alerts that relay important changes in the location, proximity, or attentional status of friends and colleagues; and,
- alerts that issue background queries based on text being composed or reviewed by users and present the results of such background searches to the user.

The invention includes computer-implemented methods, machine-readable media,
10 computerized systems, and computers of varying scopes. Other aspects, embodiments and advantages of the invention, beyond those described here, will become apparent by reading the detailed description and with reference to the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of an example computerized device in conjunction with which
15 embodiments of the invention can be practiced;

FIG. 2 is a diagram of a notification architecture according to an embodiment of the invention;

FIG. 3 is a diagram illustrating the user mechanism of the architecture of FIG. 2 in more detail, in accordance with an embodiment of the invention;

20 FIG. 4 is a diagram of a graph showing an example decay of the utility of information contained within a notification over time, according to an embodiment of the invention; and,

FIG. 5 is a flowchart of a method according to an embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanying drawings, which form a part hereof, and in which is shown by way of illustration specific exemplary embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that logical, mechanical, electrical, and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

Example Computerized Device

Referring to FIG. 1, a diagram of an example computerized device 100 in conjunction with which embodiments of the invention may be practiced is shown. The example computerized device can be, for example, a desktop computer, a laptop computer, a personal digital assistant (PDA), a cell phone, etc.; the invention is not so limited. The description of FIG. 1 is intended to provide a brief, general description of a suitable computerized device in conjunction with which the invention may be implemented. Those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PC's, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network.

The device 100 includes one or more of the following components: processor(s) 102, memory 104, storage 106, a communications component 108, input device(s) 110, a display 112, and output device(s) 114. It is noted, that for a particular instantiation of the device 100, one or more of these components may not be present. For example, a PDA
5 may not have any output device(s) 114, while a cell phone may not have storage 106, etc. Thus, the description of the device 100 is to be used as an overview as to the types of components that typically reside within such a device 100, and is not meant as a limiting or exhaustive description of such computerized devices.

The processor(s) 102 may include a single central-processing unit (CPU), or a
10 plurality of processing units, commonly referred to as a parallel processing environment. The memory 104 may include read only memory (ROM) 24 and/or random access memory (RAM) 25. The storage 106 may be any type of storage, such as fixed-media storage devices such as hard disk drives, flash or other non-volatile memory, as well as removable-media storage devices, such as tape drives, optical drives like CD-ROM's,
15 floppy disk drives, etc. The storage and their associated computer-readable media provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data. It should be appreciated by those skilled in the art that any type of computer-readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges,
20 random access memories (RAMs), read only memories (ROMs), and the like, may be used.

Because the device 100 may operate in a network environment, such as the Internet, intranets, extranets, local-area networks (LAN's), wide-area networks (WAN's),

etc., a communications component 108 can be present in or attached to the device 100.

Such a component 108 may be one or more of a network card, such as an Ethernet card, an analog modem, a cable modem, a digital subscriber loop (DSL) modem, an Integrated Services Digital Network (ISDN) adapter, etc.; the invention is not so limited.

- 5 Furthermore, the input device(s) 110 are the mechanisms by which a user indicates input to the device 100. Such device(s) 110 include keyboards, pointing devices, microphones, joysticks, game pads, satellite dishes, scanners, etc. The display 112 is how the device 100 typically shows output to the user, and can include, for example, cathode-ray tube (CRT) display devices, flat-panel display (FPD) display devices, etc. In addition, the
- 10 device 100 may indicate output to the user via other output device(s) 114, such as speakers, printers, etc.

Architecture

In this section of the detailed description, a notification architecture according to an embodiment of the invention is described, in conjunction with the diagram of FIG. 2.

- 15 The architecture 200 of FIG. 2 includes a user mechanism 202, a notification manager 204 (also referred to as an event broker), a number of notification sources 206a, 206b, . . . , 206n, and a number of notification sinks 208a, 208b, . . . , 208m. The sources are also referred to as event publishers, while the sinks are also referred to as event subscribers. There can be any number of sinks and sources. In general, the notification
- 20 manager 204 conveys notifications, which are also referred to as events or alerts, from the sources 206 to the sinks 208, based on information stored in the user mechanism 202. Each of the components of the architecture 200 of FIG. 2 is now described in turn.

The user mechanism 202 stores information regarding variables and parameters of a user that influence notification decision-making. For example, the parameters may include contextual information, such as the user's typical locations and attentional focus or activities per the time of day and the day of the week, and additional parameters conditioned on such parameters, such as the devices users tend to have in different locations. Such parameters may also be functions of observations made autonomously via one or more sensors. For example, profiles may be selected or modified based on information about a user's location as might be provided by a global positioning system (GPS) subsystem, on information about the type of device being used and/or the pattern of usage of the device, the last time a device of a particular type was accessed by the user, etc. Furthermore, as is described in more detail later in the application, automated inference may also be employed, to dynamically infer parameters such as location and attention. The profile parameters may be stored as a user profile that can be edited by the user. Beyond relying on sets of predefined profiles or dynamic inference, the notification architecture can also allow users to specify in real-time his or her state, such as the user not being available except for important notifications for the next x hours, or until a given time.

The parameters can also include default notification preference parameters regarding a user's preference as to being disturbed by notifications of different kinds in different settings, which can be used as the basis from which to make notification decisions by the notification manager 204, and to the basis upon which a particular user can make changes. The parameters may include default parameters as to how the user wishes to be notified in different situations, such as by cell phone, by pager, etc. The

parameters can include such assessments as the costs of disruption associated with being alerted by different modes in different settings. That is, in one embodiment, the parameters can include contextual parameters indicating the likelihoods that the user is in different locations, the likelihoods that different devices are available, and the likelihoods of his or her attentional status at a given time, as well as notification parameters indicating how the user desires to be notified at a given time.

The information stored by the user mechanism 202 is in one embodiment inclusive of contextual information determined by the mechanism 202. The contextual information is determined by the mechanism 202 by discerning the user's location and attentional status based on one or more contextual information sources, as is described in more detail in a later section of the detailed description. The mechanism 202, for example, may be able to determine with precision the actual location of the user via a global positioning system (GPS) that is a part of a user's car, cell phone, etc. The mechanism 202 may also use a statistical model to determine the likelihood that the user is in a given state of attention by considering background assessments and/or observations gathered through considering such information as the type of day, the time of day, the data in the user's calendar, and observations about the user's activity. This is described in more detailed in the cofiled, copending and coassigned patent application entitled "Contextual Models and Methods for Inferring Attention and Location."

[attorney docket no. 1018.101US4] The given state of attention can include whether the user is open to receiving notification, busy and not open to receiving notification, etc. The type of day can include weekdays, weekends, holidays, etc.

Each of the sources 206a, 206b, . . . , 206n is able to generate notifications intended for the user. For example, the sources 206 may include communications, such as Internet and network-based communications, local desktop computer-based communications, and telephony communications, as well as software services, such as intelligent help, background queries, automated scheduling, etc. A notification source is defined generally herein as that which generates events, which can also be referred to as notifications and alerts, intended to alert a user, or a proxy for the user, about information, services, or a system or world event. A notification source can also be referred to as an event source.

For example, e-mail may be generated as notifications by an e-mail notification source such that it is prioritized, where the host application program generating the notification assigns the e-mail with a relative priority corresponding to the likely importance or urgency of the e-mail to the user. The e-mail may also be sent without regard to their relative importance to the user. Desktop-centric notifications can include an automated dialog with the goal of alerting a user to a potentially valuable service that he or she may wish to execute (e.g., scheduling from a message), information that the user may wish to review (e.g., derived from a background query), or errors and other alerts generated by a desktop computer. Internet-related services can include notifications including information that the user has subscribed to, such as headlines of current news every so often, stock quotes, etc.

Other notifications can include background queries (e.g., while the user is working, text that the user is currently working on may be reviewed, such that background queries regarding the text are formulated and issued to search engines),

scheduling tasks from a scheduling or other program, etc. Notification sources 206 can themselves be push-type or pull-type sources. Push-type sources are those that automatically generate and send information without a corresponding request, such as headline news and other Internet-related services that send information automatically once subscribed to. Pull-type sources are those that send information in response to a request, such as e-mail being received after a mail server is polled. Still other notification sources include the following:

- e-mail desktop applications such as calendar systems;
- computer systems (e.g., that may alert the user with messages that information about alerts about system activity or problems);
- Internet-related services, appointment information, scheduling queries;
- changes in documents or numbers of certain kinds of documents in one or more shared folders;
- availability of new documents in response to standing or persistent queries for information; and,
- information sources for information about people and their presence, their change in location, their proximity (e.g., let me know when I am traveling if another employee of Microsoft is within 10 miles of me”), or their availability (e.g., let me know when Steve is available for a conversation and is near a high-speed link that can support full video conferencing”).

Each of the notification sinks 208a, 208b, . . . , 208n is able to provide the notifications to the user. For example, such notification sinks 208 can include computers, such as desktop and/or laptop computers, handheld computers, cell phones, landline

phones, pagers, automotive-based computers, etc. It is noted that some of the sinks 208 can convey notifications more richly than other of the sinks 208. For example, a desktop computer typically has speakers and a relatively large color display attached to it, as well as having a high bandwidth for receiving information when attached to a local network or to the Internet. Therefore, notifications can be conveyed by the desktop computer to the user in a relatively rich manner. Conversely, most cell phones have a small display that is black and white, and receive information at relatively low bandwidth.

Correspondingly, the information associated with notifications conveyed by the cell phone usually must be shorter and geared towards the phone's known limitations. Thus, the content of a notification may differ depending on whether it is to be sent to a cell phone or a desktop computer. A notification sink can in one embodiment refer to that which subscribes, via an event subscription service, to events or notifications.

The notification manager 204 accesses the information stored by the user mechanism 202, and makes a decision as to which of the notifications it receives from the sources 206 to convey to which of the sinks 208 based on this information. Furthermore, the manager 204 is able to determine how the notification is to be conveyed, depending on which of the sinks 208 it has selected to send the information to. For example, it may determine that the notification should be summarized before being provided to a given of the sinks 208. The manager 204 is in one embodiment a computer program executed by a processor of a computer from a machine-readable medium such as a memory thereof.

The invention is not limited to how the manager 204 makes its decisions as to which of the notifications to convey to which of the notification sinks, and in what manner the notifications are conveyed. In one embodiment, a decision-theoretic analysis

can be made. For example, the notification manager can be designed to infer important uncertainties about variables including a user's location, attention, device availability, and amount of time until the user will access the information if there were no alert. The notification manager can then make notification decisions about whether to alert a user to a notification, and if so, the nature of the summarization and the best device or devices to employ for relaying the notification. Such a particular embodiment is described in the cofiled, copending and coassigned case entitled "Decision-Theoretic Principles and Policies for Notification." [attorney docket no. 1018.101US3] In general, the manager determines the net expected value of a notification. In doing so, it can consider the following:

- the fidelity and transmission reliability of each available notification sink;
- the attentional cost of disturbing the user;
- the novelty of the information to the user;
- the time until the user will review the information on his or her own;
- the potentially context-sensitive value of the information; and,
- the increasing and/or decreasing value over time of the information contained within the notification.

The inferences made about uncertainties thus may be generated as expected likelihoods of values such as the cost of disruption to the user with the use of a particular mode of a particular device given some attentional state of the user, etc. The notification manager 204 makes decisions as to one or more of the following:

- what the user is currently attending to and doing (based on, for example, contextual information);

- where the user currently is;
 - how important the information is;
 - what is the cost of deferring the notification;
 - how distracting would a notification be;
- 5 • what is the likelihood of getting through to the user; and,
- what is the fidelity loss associated with the use of a specific mode of a given notification sink.

Therefore, ultimately, the notification manager 204 performs an analysis, such as a decision-theoretic analysis, of pending and active notifications, evaluates context-

10 dependent variables provided by information sinks and sources, and infers key uncertainties, such as the time until a user is likely to review provided information and the user's location and current attentional state.

As used herein, inference refers generally to the process of reasoning about or inferring states of the system, environment, or user from a set of observations as captured

15 via events and/or data. Inference can be used to identify a specific context or action, or can be employed to generate a probability distribution over states. The inference can be probabilistic – that is, the computation of a probability distribution over states of interest based on a consideration of data and events. Inference can also refer to techniques used for composing higher-level events from a set of events and/or data. Such inference

20 results in the construction of new events or actions from a set of observed events and/or stored event data, whether or not the events are correlated in close temporal proximity, and whether the events and data come from one or several event and data sources.

Furthermore, in one embodiment, the notification manager 204 accesses information stored in a user profile by the user mechanism 202 in lieu of or to support a personalized decision-theoretic analysis. For example, the user profile may indicate that at a given time, the user prefers to be notified via a pager, and only if the notification has a predetermined importance level. Such information can be used as a baseline from which to start a decision-theoretic analysis, or can be the only manner by which the manager 204 determines how and whether to notify the user.

In one embodiment of the invention, the notification platform architecture is constructed as a layer that resides over an eventing or messaging infrastructure.

However, the invention is not limited to any particular eventing infrastructure. Such eventing and messaging systems and protocols include:

- HyperText Transport Protocol (HTTP), or HTTP extensions as known within the art;
- Simple Object Access Protocol (SOAP), as known within the art;
- Windows Management Instrumentation (WMI), as known within the art;
- Jini, as known within the art; and,
- any type of communications protocols, such as those based on packet-switching protocols, etc.

Furthermore, in one embodiment, the architecture is constructed as a layer that resides over a flexible distributed computational infrastructure, as can be appreciated by those of ordinary skill within the art. Thus, the notification platform architecture can utilize an underlying infrastructure as a manner by which sources send notifications, alerts and

events, as a manner by which sinks receive notifications, alerts and events, etc. The invention is not so limited, however.

In the following sections of the detailed description, the user mechanism 202, the notification sources 206, and the notification sinks 208 are each described in more detail, according to varying embodiments of the invention. Thereafter, a method of the manner by which the architecture operates according to an embodiment of the invention is presented.

User Mechanism

In this section of the detailed description, the user mechanism of the notification architecture described in the previous section of the detailed description is described in more detail, in conjunction with FIG. 3, which is a diagram 300 showing the user mechanism 202 of FIG. 2 in more detail. The user mechanism 202 as shown in FIG. 3 includes a user notification preferences store 302, a user context module 304 that includes a user context profile store 305, and a whiteboard 307. The user mechanism 202 is in one embodiment implemented as one or more computer programs executable by a processor of a computer from a machine-readable medium thereof, such as a memory.

The store 302 stores notification parameters for a user, such as the default notification preferences for the user, as a user profile, which can then be edited and modified by the user. The store 302 can be considered as that which stores information on parameters that influence how a user is to be notified. The user context module 304 determines a user's current context, based on the context information sources 306 as published to the whiteboard 307; the user context profile store 305 stores the context parameters for a user, such as the default context settings for the user, which can be

edited and modified by the user. That is, the user context module 304 provides a best guess about current context information by accessing information from the store 305 and/or updating the prior set of beliefs in the store 305 with live sensing, via one or more context sources 306. The store 305 can be considered as that which stores a priori where
5 a user is, and what the user is doing. The module 304 is in one embodiment implemented as a computer program executable by a processor of a computer from a machine-readable medium thereof, such as a memory.

The user context profile store 305 is a pre-assessed or predefined user profile that captures such information as a deterministic or probabilistic profile. The profile can be
10 of typical locations, activities, device availabilities, and costs and values of different classes of notification as a function of such observations as time of day, type of day, and user interactions with one or more devices. The type of day can include weekdays, weekends, holidays, etc. The user context module 304 can then actively determine or infer key aspects of the context of the user, such as the user's current or future location
15 and attentional state. Furthermore, the actual states of context can be accessed directly from the sources 306 via the whiteboard 307, or, can be inferred from a variety of such observations through inferential methods such as Bayesian reasoning as known within the art.

The context information sources 306 provide information to the module 304 via
20 the whiteboard 307 regarding the user's attentional state and location, from which the module 304 can make a determination as to the user's current context (that is, the user's current attentional state and location). Furthermore, the invention is not limited to a particular number or type of context sources 306, nor the type of information inferred or

accessed by the user context module. However, in one embodiment, the context sources 306 include multiple desktop information and events, such as mouse information, keyboard information, application information (e.g., which application is currently receiving the focus of the user), ambient sound and utterance information, text information in the windows on the desktop, etc. The whiteboard 307 is in one embodiment a common storage area, to which the context information sources 306 can publish information, and from which multiple components, including sources and the module 304 can access this information; it is not necessary for implementation of the invention, however. An event, also referred to as a notification or alert, can in one embodiment generally include information about an observation about one or more states of the world. Such states can include the status of system components, the activity of a user, or a measurement about the environment. Furthermore, events can be generated by an active polling of a measuring device or source of events, by the receipt of information that is sent on a change, or per a constant or varying event heartbeat.

Other types of context sources 306 include personal-information manager (PIM) information of the user, which usually can provide scheduling information regarding the schedule of the user, etc. The current time of day, as well as the user's location – for example, determined by a global positioning system (GPS), or a user's use of a cell phone, PDA, or a laptop that can be locationally determined – are also types of context sources 306. Furthermore, real-time mobile device usage is a type of context source 306. For example, a mobile device such as a cell phone may be able to determine if it is currently being used by the user, as well as its orientation and tilt (indicating information

regarding its usage as well), and acceleration and speed (indicating information as to whether the user is moving or not).

Notification Sources

In this section of the detailed description, the notification sources 206 of the architecture 200 of FIG. 2 are described in more detail. The notification sources generally generate notifications that are conveyed to the notification manager, which determines when notifications should occur, and, if so, which of the notifications should be conveyed to which of the notification sinks and in what order, as has also been already described.

In one embodiment, each notification source has identified with it one or more of the following parameters within a standard description of key attributes and relationships, referred to herein as a notification source schema or source schema. It is noted that schema are provided for sources, for sinks, and for context-information sources. Such schemas provide declarative information about different component and provide a mechanism for the sources, the notification manager, the sinks, and the user mechanism (including the context mechanism described in the previous section of the detailed description) to share semantic information with one another. Thus, different schemas provide information about the nature, urgency, and device signaling modalities associated with notification. That is, schema can be defined generally as a collection of classes and relationships among classes that defines the structure of notifications and events, containing information including event or notification class, source, target, event or notification semantics, ontological content information, observational reliability, and any quality-of-service attributes.

The parameters for each notification source schema include one or more of:
message class; relevance; importance; time criticality; novelty; content attributes; and
fidelity tradeoffs, and source information summary information. The message class for a
notification generated by a notification source indicates the type of communication of the
notification, such as e-mail, instant message, numerical financial update, desktop service,
etc. The relevance for a notification generated by a notification sources indicates a
likelihood that the information contained within the notification is relevant, for one or
more specified contexts. In one embodiment, the relevance is a logical flag, indicating
whether the source is relevant for a given context or not. The novelty of the notification
indicates the likelihood that the user already knows the information contained within the
notification. That is, the novelty is whether the information is new to the user, over time
(indicating if the user knows the information now, and when, if ever, the user will learn
the information in the future without being alerted to it).

Fidelity tradeoffs associated with the notification indicate the loss of value of the
information within the notification that can result from different forms of specified
allowed truncation, summarization, etc. Such truncation and/or summarization may be
required for the notification to be conveyed to certain types of notification sinks that may
have bandwidth or other limitations preventing them from receiving the full notification
as originally generated. Fidelity in general refers to the nature and/or degree of
completeness of the original content associated with a notification. For example, a long
e-mail message may be truncated, or otherwise summarized to the maximum of 100
characters allowed by a cell phone, incurring a loss of fidelity. Likewise, an original
message containing text and graphics content suffers a loss in fidelity when transmitted

via a device that only has text capabilities. In addition, a device may only be able to show a portion of the full resolution available from the source. Fidelity tradeoffs refer to a set of fidelity preferences of a source stated either in terms of orderings (e.g., rendering importance in order of graphics first, then sound) and/or costs functions that indicate how the total value of the content of the notification diminishes with changes in fidelity. For example, a fidelity tradeoff might describe how the full value associated with the transmission of a complete e-mail message changes with increasingly greater amounts of truncation. Content attributes contain a summary of the nature of the content, representing such information as whether the core message includes text, graphics, and audio components. The content itself is the actual graphics, text, and/or audio that make up the message content of the notification.

The importance of a notification refers to the value of the information contained in the notification to the user, assuming the information is relevant in the current context. In one embodiment, it is expressed as a dollar value of the information's worth to the user. Time criticality indicates the time-dependent change in the value of the information contained in the notification – that is, how the value of the information changes over time. In most but not all cases, the value of the information of a notification decays with time. This is particularly shown in the diagram 400 of FIG. 4. The graph 402 shows the utility of a notification mapped over time. At the point 404 within the graph, representing the initial time, the importance of the notification is indicated, while the curve 406 indicates the decay of the utility over time.

Default attributes and schema templates for different notification sources or source types may be made available in notification source profiles stored in the user

notification preferences store, such as the store 302 of FIG. 3 as described in the previous section. Such default templates can be directed to overrule values provided by notification sources or to provide attributes when they are missing from schema provided by sources. Source summary information allows a source to post general summaries of the status of information and potential notifications available from a source. For example, source summary information from a messaging source might include information about the total number of unread messages that are at least some priority, the status of attempts by people to communicate with a user, etc.

Notification Sinks

In this section of the detailed description, the notification sinks 208 of the architecture 200 of FIG. 2 are described in more detail. The notification sinks are devices, etc., by which the user can be notified of information contained in notifications. The choice as to which sink or sinks are to be used to convey a particular notification is determined by the notification manager.

In one embodiment, each notification sink has identified with it one or more of the following parameters within a schema: device class; modes of signaling (alerting); and, for each mode, fidelity/rendering capabilities, transmission reliability, actual cost of communication, and attentional cost of disruption. For devices that allow for the parameterized control of alerting attributes, the schema for the devices can additionally include a description of the alerting attributes and parameters for controlling the attributes, and functions by which other attributes (e.g., transmission reliability, cost of distribution, etc.) change with the different settings of the alerting attributes. The schema for notification sinks provides for the manner by which the notification devices

communicate semantic information about their nature and capabilities with the notification manager and other components of the system. Default attributes and schema templates for different device types can be made available in device profiles stored in the user notification preferences store, such as the store 302 of FIG. 3 as described in the

5 previous section. Such default templates can be directed to overrule values provided by devices or to provide attributes when they are missing from schema provided by devices.

Each of the schema parameters is now described in term. The class of the device refers to the type of the device: a cell phone, a desktop computer, a laptop computer, etc. The class can also be more general, such as a mobile device, a stationery device, etc. The

10 modes of signaling refer to the different ways in which a given device can alert the user about a notification. Devices may have one or more notification modes. For example, a cell phone can only vibrate, only ring with some volume, or it can both vibrate and ring. Furthermore, a desktop display for an alerting system can in one embodiment be decomposed into several discrete modes (e.g., a small notification window in the upper

15 right hand of the display vs. a small thumbnail at the top of the screen – with or without an audio herald), as described in the cofiled, copending and coassigned application entitled “Display and Human-Computer Interaction for a Notification Platform.”

[attorney docket no. 1018.101US2] Beyond being limited to a set of predefined behaviors, a device can allow for modes with alerting attributes that are functions of

20 parameters, as part of its definition. Such continuous alerting parameters for a mode represent such controls as the volume at which an alert is played at the desktop, rings on a cell phone, the size of an alerting window, etc.

limit may be 100 characters for any single message, and the phone may have no graphics abilities.

Methods

In this section of the detailed description, methods according to varying
5 embodiments of the invention are shown. The methods can in some embodiments be
computer-implemented. A computer-implemented method is desirably realized at least in
part as one or more programs running on a computer -- that is, as a program executed
from a computer-readable medium such as a memory by a processor of a computer. The
programs are desirably storable on a machine-readable medium such as a floppy disk or a
10 CD-ROM, for distribution and installation and execution on another computer. The
program or programs can be a part of a computer system or a computer, such as that
described in conjunction with FIG. 1 in a previous section of the detailed description.
The invention is not so limited, however.

Referring to FIG. 5, a flowchart of a method 500 is shown. In one embodiment,
15 the method 500 can be executed by the notification architecture 200 of FIG. 2, which has
been already described. In 502, one or more notification sources generate notifications,
which are received by a notification manager. In 504, the user mechanism generates
context information regarding the user, which in 506 is received by the notification
manager. That is, in one embodiment, in 504, the user mechanism accesses a user
20 contextual information profile that indicates the user's current attentional status and
location, and/or assesses real-time information regarding the user's current attentional
status and location from one or more contextual information sources, as has been
described in the previous sections of the detailed description.

In 508, the notification manager determines which of the notifications to convey to which of the notification sinks, based in part on the context information received from the user mechanism. The notification also makes its determination based on information regarding notification parameters of the user as stored by the user mechanism, in one embodiment. That is, in one embodiment, in 508, the manager performs a decision-theoretic analysis as to whether a user should be alerted for a given notification, and how the user should be notified. Notification parameters regarding the user can be utilized to personalize the analysis by filling in missing values or by overwriting parameters provided in the schema of sources or sinks. Notification preferences can also provide policies that are employed in lieu of the decision-theoretic analysis, as has been described in the previous sections of the detailed description. Based on this determination, the notification manager conveys the notifications to the sinks in 510.

Application of the Invention to Entities Other than Users

Embodiments of the invention have been described herein thus far as applicable to users. However, the invention itself is not so limited. That is, the invention is applicable to any type of entity, including users. Other types of entities include agents, processes, computer programs, threads, services, servers, computers, machines, companies, organizations, businesses, etc. The agent, for example, may be a software agent, which can be generally defined as a computer program that performs a background task for a user and reports to the user when the task is done or some expected event has taken place. Still other types of entities are encompassed under the invention, as can be appreciated by those of ordinary skill within the art.

Thus, where embodiments of the invention have been particularly described in relation to a user, the invention itself is applicable to any type of entity. For example, the user mechanism of a system according to an embodiment of the invention can be generalized as a mechanism, applicable to any type of entity. As another example,
5 notification sinks can generate notifications, alerts and events regarding and for entities other than users. Similarly, notification sinks can receive notifications, alerts and events regarding and for entities other than users.

Conclusion

It is noted that, although specific embodiments have been illustrated and
10 described herein, it will be appreciated by those of ordinary skill in the art that any arrangement that is calculated to achieve the same purpose may be substituted for the specific embodiments shown. This application is intended to cover any adaptations or variations of the present invention. Therefore, it is manifestly intended that this invention be limited only by the claims and equivalents thereof.